QoS Mechanisms and OVS for NFV

Al Morton

7 March 2016

Outline

- Background
- Two Class Scenario
- Observations so far
- Solutions may involve new mechanisms
- Possible Next Steps

Traditional QoS mechanisms and Example of QoS Classes

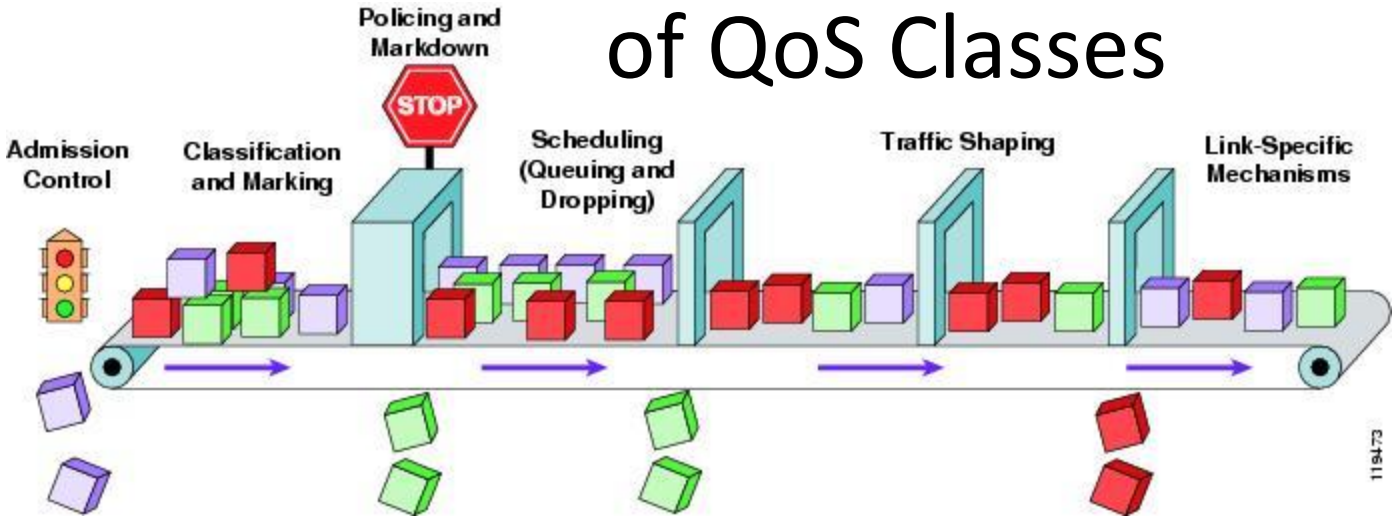

Note: If the resources are sufficient to handle all traffic, QoS Mechanisms take little/no action.

QCI	Bearer	Priority	Delay	PELR	Examples
1	GBR	2	100 ms	10^{-2}	Conversational voice
2		4	150 ms	10^{-3}	Conversational video
3		3	50 ms	10^{-3}	Real-time games
4		5	300 ms	10^{-6}	Streaming video
5	Non-GBR	1	100 ms	10^{-6}	IMS signalling
6		6	300 ms	10^{-6}	Streaming video, web, EMail
7		7	100 ms	10^{-3}	Voice, video, games
8		8	300 ms	10^{-6}	Streaming video, web, EMail
9		9			

Traditional QoS mechanisms and Example of QoS Classes

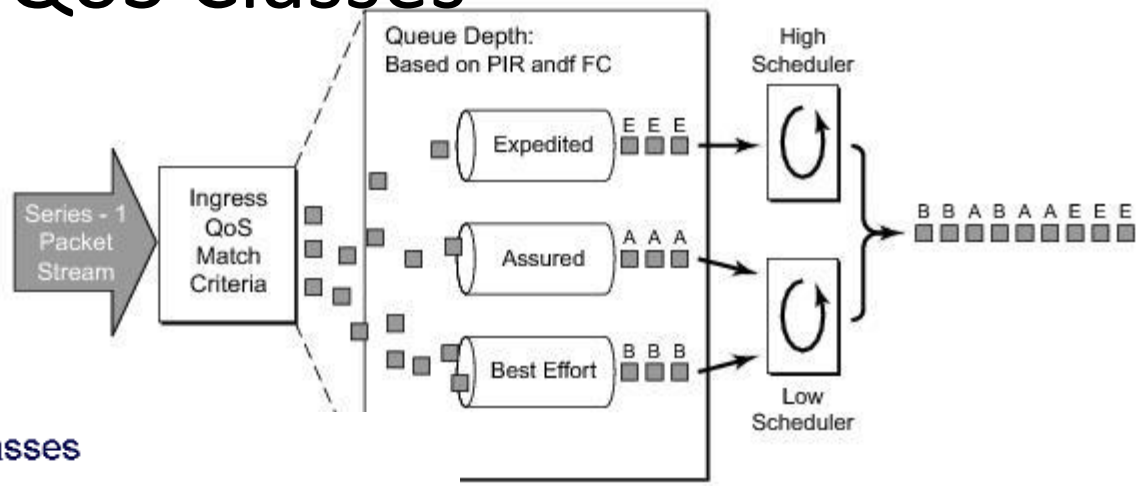

ITU-T Y.1541 QoS Classes

Network Performance Parameter	Y.1541 QoS Classes					
	Class 0	Class 1	Class 2	Class 3	Class 4	Class 5
IPTD	100ms	400ms	100ms	400ms	1 s	U
IPDV	50ms	50ms	U	U	U	U
IPLR	$1*10^{-3}$	$1*10^{-3}$	$1*10^{-3}$	$1*10^{-3}$	$1*10^{-3}$	U
IPER	$1*10^{-4}$					U

Classes resolve scaling issues

Example Service Mappings

- Voice
- Signaling
- Interactive Data
- Best Effort Data
- Streaming Video

Note that delay variability is the distinguishing difference between

Two Class Scenario

- “Real-time” Class requires priority treatment
  - But not Strict priority, resources must be shared
- “Default Data” Class can experience less performance, more impairments
- Each Class is reliably “marked” in one or more packet/frame headers. Expect variable 4-tuple
- Encapsulation/Tunnels should expose the markings (possibly mapped to alt code)
- Traffic may originate/terminate on NIC or VM

Observations so far…

- Support for packet classification *early* in the processing pipeline
- Queue-based mechanisms have limited scope
  - OVS processes each packet to completion (?)
  - Queue between vSwitch and Guest/VM
- Assuming OpenFlow Control, possible to mark some flows with higher priority
  - But OVS “flattens” many rules into 1 in fastpath
  - Possibly requires support in ODL and OS, too
- Does testing indicate there is a (treatable) problem with the 2 class scenario?

Solutions may involve new mechanisms: OF control

- OF 1.3+ adds Auxiliary Control Connection
  - Parallel to Main Connection (sec 6.3.6)
- For each of two classes, add
  - Aux 1 connection for “Real-time” PACKET_IN
  - Aux 2 connection for “Default” PACKET_IN
- Switch may process Aux 1 responses with higher priority than Aux 2
- Controller *could* offer similar policy
  - Upstream improvement in ODL, etc.
- But, need *benefits* of such a mechanism

Possible Next Step

- Tests with OVS, OVS+DPDK
  - Use multiple flows in each class
  - Classes have different packet sizes consistent with traditional uses:
    - “Real-time” has small (~240 byte) packets at 50 pkt/s
    - “Default Data” has max MTU packets in one direction, and TCP ACKs in the other. Streams tend to be bursty with many MTU size packets back-to-back.
  - Check for delay and delay variation on the Real-time class.
  - Note: Bottleneck project tests showed an odd throughput limitation vs packet size – needs some more investigation.

Work further at Hackfest?

- Al is attending…